U.S. HOUSE OF REPRESENTATIVES
COMMITTEE ON SCIENCE

HEARING CHARTER

H.R. 4218, the High-Performance Computing Revitalization Act of 2004

Thursday, May 13, 2004
10:30 a.m. - 12:30 p.m.
2318 Rayburn House Office Building

1. Purpose

On Thursday, May 13, 2004, the House Science Committee will hold a hearing to examine federal high-performance computing research and development (R&D) activities and to consider H.R. 4218, the *High-Performance Computing Revitalization Act of 2004*, which would amend the *High-Performance Computing Act of 1991*.

The bill is timely because high-performance computing in the U.S. is at a turning point. The fastest computer in the world today is in Japan not the U.S., and several federal agencies are in the process of reformulating their high-performance computing programs, in part, in response to the challenge posed by Japan.

2. Witnesses

Dr. John H. Marburger III is the Director of the White House Office of Science and Technology Policy (OSTP). Prior to joining OSTP, Dr. Marburger served as President of the State University of New York at Stony Brook and as Director of the Brookhaven National Laboratory.

Dr. Irving Wladawsky-Berger is Vice President for Technology and Strategy for IBM Corporation. Dr. Wladawsky-Berger previously served as co-chair of the President's Information Technology Advisory Committee (PITAC), and as a founding member of the Computer Sciences and Telecommunications Board of the National Academy of Sciences.

Dr. Rick Stevens is the Director of the Mathematics and Computer Science Division at Argonne National Laboratory (ANL). He is also a Director of the National Science Foundation (NSF) TeraGrid project, which aims to build the nation's most comprehensive, open infrastructure for scientific computing.

Dr. Daniel Reed is the William R. Kenan, Jr. Eminent Professor at the University of North Carolina at Chapel Hill (UNC-CH). Previously, Dr. Reed served as Director of the National Center for Supercomputing Applications at the University of Illinois Urbana-Champaign, one of NSF's university-based centers for high-performance computing. Dr. Reed is a current member of PITAC.

3. Overarching Questions

The hearing will address the following overarching questions:

1. How does high-performance computing affect the international competitiveness of the U.S. scientific enterprise?

2. Are current efforts on the part of the federal civilian science agencies in high-performance computing sufficient to assure U.S. leadership in this area? What should agencies such as the NSF and the Department of Energy (DOE) be doing that they are currently are not?

3. Where should the U.S. be targeting its high-performance computing research efforts? Are there particular industrial sectors or science and engineering disciplines that will benefit in the near term from anticipated high-performance computing developments?

4. Brief Overview

- High-performance computers (also called supercomputers or high-end computers) are an essential component of U.S. scientific, industrial, and military competitiveness. However, the fastest and most efficient supercomputer in the world today is in Japan, not the U.S. Japan was successful in producing a computer far ahead of the American machines in part because Japan focused on a type of computer architecture that the U.S. had ceased developing. Also, Japan focused a large amount of money on a single machine, while the U.S. funds a variety of computer development projects.

- Despite the recent technical success of the Japanese, most experts still rate the U.S. as highly competitive in high-performance computing. The depth and strength of U.S. capability stems in part from the sustained research and development program carried out by federal science agencies under an interagency program codified by the *High-Performance Computing Act of 1991*. That Act is widely credited with reinvigorating U.S. high-performance computing capabilities after a period of relative decline during the late 1980s.

- The Federal government promotes high-performance computing in several different ways. First, it funds research and development (R&D) at universities, government laboratories and companies to help develop new computer hardware and software; second, it funds the purchase of high-performance computers for universities and government laboratories; and third, it provides access to high-performance computers for a wide variety of researchers by allowing them to use government-supported computers at universities and government labs.

- In recent years, federal agency efforts once again appear to have lost momentum as federal computing activities began focusing less on high-performance computing and more on less specialized computing and networking technologies.

- Responding to concerns that U.S. efforts to develop and deploy high-performance computers may have flagged, OSTP created an interagency task force—the High-End Computing

Revitalization Task Force (HEC-RTF)—to examine federal high-performance computing programs and make recommendations for improvement. Dr. Marburger will release the task force report during his appearance before the Committee.

- On April 27, 2004, Representative Judy Biggert introduced H.R. 4218, the *High-Performance Computing Revitalization Act of 2004*, which would update the *High-Performance Computing Act of 1991* and, in particular, would require the High-Performance Computing R&D Program to "provide for sustained access by the research community in the United States to high-performance computing systems that are among the most advanced in the world in terms of performance in solving scientific and engineering problems, including provision for technical support for users of such systems." H.R. 4218 also requires the Director of OSTP to "develop and maintain a research, development, and deployment roadmap for the provision of high-performance computing systems for use by the research community in the United States." This and other provisions in the bill are designed to ensure a robust ongoing planning and coordination process so that the national high-performance computing effort is not allowed to lag in the future.

5. Major Issues Addressed in H.R. 4218

Assuring U.S. Researchers Access to the Most Advanced High-Performance Computing Systems Available.

What the Bill Does: The bill requires the High-Performance Computing Research and Development Program to "provide sustained access by the research community in the United States to high-performance computing systems that are among the most advanced in the world in terms of performance in solving scientific and engineering problems, including provision for technical support for users of such systems." The bill also specifically requires the NSF and the DOE Office of Science to provide U.S. researchers with access to "world class" high-performance computing systems.

Why that's necessary: Beginning in the 1980s with the NSF supercomputer centers program, the federal government has been providing university researchers with access to the fastest computers. Today, university researchers are concerned that the federal government, and particularly NSF, may be moving away from a commitment to provide such access. While NSF has reiterated its intention to continue to provide access to the fastest computers through supercomputer centers, it has also said it will place greater emphasis on distributed collections of many computers (known as "grid computing"), which may not provide computing capability equal to that of the fastest supercomputers. At the same time, DOE has indicated it wants to expand its efforts to provide access to large, single-location machines, but it is not clear how much access DOE will be able to provide or whether its machines will be open to researchers in all fields as NSF-funded machines are.

Assuring Balanced Progress on All Aspects of High-Performance Computing.

What the Bill Does: The bill also requires the program to support all aspects of high-performance computing for scientific and engineering applications, including software, algorithm and applications development, development of technical standards, development of new computer models for science and engineering problem solving, and education and training in all the disciplines that support advanced computing.

Why that's necessary: New supercomputers (hardware) alone won't help researchers. The development of advanced software and applications programs is essential to enable researchers to use the additional computing power.

Assuring an Adequate Interagency Planning Process to Maintain Continued U.S. Leadership.

What the Bill Does: The bill requires the Director of OSTP to "develop and maintain a research, development, and deployment roadmap for the provision of high-performance computing systems for use by the research community in the United States." This and other provisions in the bill are designed to ensure a robust ongoing planning and coordination process so that the national high-performance computing effort is not allowed to lag in the future.

Why that's necessary: The High-Performance Computing Act of 1991 codified an interagency planning process that remains in place today. However, the chief product of this process in recent years has been an annual review of activities undertaken by agencies, rather than a prospective planning document. A forward-looking process would enhance coordination between agencies and maximize the total benefit of federal investment.

6. Current Issues in High-Performance Computing

Is the U.S. Competitive?

The world's fastest computer, Japan's Earth Simulator, is designed to perform simulations of the global environment and to address scientific questions related to climate, weather, and earthquakes. NEC, a leading Japanese computer manufacturer, built the Earth Simulator for the Japanese government at a cost of at least $350 million. The first measures of the Earth Simulator's speed, taken in April 2002, determined that the Earth Simulator was significantly faster than the former record holder—the ASCI White System at Lawrence Livermore National Laboratory—and also used the machine's computing power with far greater efficiency.[1]

Twice a year, researchers at the University of Tennessee and the University of Mannheim (United Kingdom) compile a list of the world's 500 fastest supercomputers. The latest list became public on November 16, 2003 (see Table 2 in Appendix II).[2] The Earth Simulator is approximately twice as fast as the second place machine, the ASCI Q system (located at Los

---

[1] For the fastest U.S. computers, typical scientific applications are usually only able to utilize 5-10 percent of the theoretical maximum computing power, while the design of the Earth Simulator makes 30-50 percent of its power accessible to the majority of typical scientific applications.

[2] The top 500 list is compiled by researchers at the University of Mannheim (Germany), Lawrence Berkeley National Laboratory, and the University of Tennessee and is available on line at http://www.top500.org/. For a machine to be included on this public list, its owners must send information about its configuration and performance to the list-keepers. Therefore, the list is not an entirely comprehensive picture of the high-performance computing world, as classified machines, such as those used by NSA, are not included.

Alamos National Laboratory and built by Hewlett-Packard). Of the top twenty machines, eight are located at DOE national laboratories and two at U.S. universities.[3] IBM manufactured six of the top twenty machines and Hewlett-Packard manufactured five.

What Types of High-Performance Computers Should the U.S. Develop?

The success of the Earth Simulator has caused a great deal of soul-searching in the high-performance computing community in the U.S. The Earth Simulator is built from custom-made components, and is based on a computer architecture that the U.S. had stopped pursuing in the 1990s. At that time, U.S. programs chose to favor the use of commercially available components for constructing high-performance computers. An advantage of this approach was that it made high-performance computers more cost-effective to develop, by leveraging development costs against a larger market.

Some computing experts have concluded that this strategy of relying largely on commercial needs to guide the development of supercomputer components has left U.S. academic researchers at a disadvantage. That's because certain kinds of research questions – such as those involved in climate modeling – are difficult to pursue on the kinds of computers that can be built with commercial components. The Japanese Earth Simulator, for example, is not based on a computer architecture that would be of widespread interest in the commercial market.

Federal agencies are in the process of reviewing their programs to decide which kinds of computer architecture R&D to pursue. H.R. 4218 is silent on this issue, but a decision on what kinds of computer architectures to pursue would be part of the planning required by the bill.

This question is significant in that NSF first became involved in offering supercomputer access because in the early 1980s foreign researchers often had more and better access to top supercomputers than U.S. researchers did. With the advent of the Earth Simulator, this may be true again for climate and earthquake researchers. Federal civilian agencies, particularly NSF, need to figure out how to help develop computers that will be useful to U.S. scientists in a wide variety of fields. The research needs of different scientific fields require distinct computer architectures, and so serving the entire user community will most likely require the development of a number of diverse computer architectures.

Supercomputers – regardless of the extent of their appeal in the commercial market – are still in the end manufactured private companies. In the U.S., the major producers of high-performance computers include IBM, Hewlett-Packard, and Silicon Graphics, Inc. and Cray. Leading Japanese manufacturers include NEC, Fujitsu, and Hitachi. In the past, Congress prevented federal research funds from being used to purchase Japanese supercomputers.

---

[3] The two university machines are located at the Pittsburgh Supercomputing Center (supported primarily by NSF) and Louisiana State University's Center for Applied Information Technology and Learning. The remaining 12 machines include four in Europe, two in Japan, and one each at the National Oceanic & Atmospheric Administration, the National Center for Atmospheric Research, the Naval Oceanographic Office, and NASA.

Where are the NSF and DOE Office of Science and Programs Headed?

NSF and the DOE Office of Science are the lead agencies responsible for providing high-performance computing resources for U.S. civilian research. (See Appendix II.) Both NSF and the DOE Office of Science are moving ahead in significant new directions. NSF recently signaled that it will place greater emphasis on developing grid computing resources. Meanwhile, DOE has indicated it will expand its efforts to provide access to large, single-location machines but has not yet implemented these plans. Both agencies are at a point of transition as they redefine their roles in providing access to U.S. researchers to high-performance computing resources.

NSF's support three large supercomputer centers,[4] which in FY03 served approximately 3,000 users, mostly from academia. (When the supercomputer center program started, there were five initial centers.) In addition to providing cyberinfrastructure, NSF's Computer and Information Sciences and Engineering Directorate supports roughly $70 million of research on hardware, systems architecture, and advanced applications.

In FY04, the DOE Office of Science initiated a new effort in the development of next-generation computer architectures (NGA). The program will emphasize the development of computer architectures that do not rely on commercial components or computing needs. The Department issued an initial request for proposals for the NGA program in March 2004. The NGA Program received $38 million in FY04, and the same amount is requested for FY05.

DOE also administers the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory, which provides high-end computing resources to over 2,000 scientists annually. According to Department figures, 35 percent of NERSC users are university-based, but the majority are those are funded through DOE grants. The budget for NERSC is on an upward trend, up from $22 million in FY03 to $32 million in FY04, with $38 million proposed for FY05. These increases reflect the Office of Science strategy to expand its role as a provider of high-performance computing resources.

Also, NSF and the Defense Advanced Research Projects Agency (DARPA) have jointly released a solicitation for software for high-performance computing (NSF/DARPA).[5]

7. Background

What is High-Performance Computing?

High-performance computing—also called supercomputing, high-end computing, and sometimes advanced scientific computing—refers to the use of machines or groups of machines that can perform very complex computations very quickly. High-performance computers are, by definition, the most powerful computers in the world at a given moment in time. High-

[4] The three NSF-supported centers are the San Diego Supercomputing Center at the University of California-San Diego, the National Center for Supercomputing Applications at the University of Illinois Urbana-Champaign, and the Pittsburgh Supercomputing Center, jointly run by Carnegie Mellon University and the University of Pittsburgh.
[5] The NSF/DARPA solicitation for research on software and tools for high-end computing is available on line at http://www.nsf.gov/pubs/2004/nsf04569/nsf04569.htm.

performance computers are used to solve highly complex scientific and engineering problems, or to manage vast amounts of data. Technologies improve so quickly that the high-performance computing achievements of a few years ago could now be handled by today's desktops.

The speed of high-performance computers is measured in "flops," a unit signifying a calculation each second. The prefix "Tera" signifies trillions, and thus a one Teraflop machine can execute a trillion calculations each second. The world's fastest machine, Japan's Earth Simulator, can execute 35 Teraflops, or 35 trillion calculations each second.

What is High-Performance Computing Used For?

High-performance computers are often used to simulate physical systems that are difficult to study experimentally. Such simulations can be an alternative to actual experiments (e.g. for nuclear weapon testing and climate modeling), or can test researchers' understanding of a system (e.g. for particle physics and astrophysics). Industry researchers use high-performance computers to simulate how new products will behave in different environments (e.g. for development of new industrial materials). Other major uses for supercomputers include performing massive mathematical calculations (e.g. for codebreaking) and managing vast amounts of data (e.g. for government personnel databases).

> *Scientific Applications:* High-performance computers are used to tackle a rich variety of scientific problems. Large-scale climate modeling examines possible future scenarios related to global warming. In biology and biomedical sciences, researchers perform simulations of protein structure and folding, and also model blood flows. Astrophysicists model planet formation and supernova, while cosmologists simulate conditions in the early universe. Particle physicists perform complex calculations involving the basic building blocks of matter. Geologists model stresses within the earth to study plate tectonics, while civil engineers simulate the impact of earthquakes.

> *National Defense Applications:* The National Security Agency (NSA) is a major user and developer of high-performance computers for specialized tasks relevant to codebreaking (such as factoring large numbers). The DOE National Nuclear Security Administration (NNSA) is also a major user and developer of machines used in modeling nuclear weapons. The Department of Homeland Security uses high-performance computing to extract useful data from large amounts of information; to model the dispersal of plumes of biological, chemical, and radiological agents; and to identify pathogens using their DNA signatures. The Department of Defense uses high-performance computing to model armor penetration, and for weather forecasting. Many scientific applications may have future defense applications. For example, computational fluid dynamics studies could be used to model turbulence surrounding military aircraft.

> *Industrial Applications:* The automotive industry uses high-performance computers for vehicle design and engineering. The movie industry uses massive computer animation programs to produce films. Pharmaceutical companies simulate chemical interactions to design new drugs. The commercial satellite industry manages huge amounts of data in generating maps. Financial companies and other industries use large computers to process immense and unpredictable Web transaction volumes, to mine databases for sales patterns or fraud, and to measure the risk in investment portfolios.

What Types of High-Performance Computers Are There?

There are a number of different ways to build high-performance computers, and different configurations are better suited to different problems. While there are many possible configurations, they can be roughly divided into two classes: big, single-location machines and distributed collections of many computers (this approach is often called grid computing). Each approach has its benefits—the big machines can be designed for a specific problem and are often faster, while grid computing is attractive in part because the purchase and storage cost is often lower than for a large specialized supercomputer.

At least since the mid-1990's, the U.S. approach to developing new capabilities has emphasized using commercially-available components as much as possible. This emphasis has resulted in an increased focus on grid computing, and has influenced the designs of large, single-location machines. The U.S. has favored supercomputer designs based on ever-larger numbers of commercially available processors, coupled with improvements in information sharing between processors.

Users thus have a number of options for high-performance computing, and must take into account the pros and cons of different configurations when deciding what sort of machine to use. Users must also design software to allow the machine to solve each problem most efficiently. For example, some problems, such as climate modeling and codebreaking, require a great deal of communication between computer components. Other applications, such as large-scale data analysis for high energy physics experiments or bioinformatics projects, can be more efficiently performed on distributed machines, each tackling its own piece of the problem in relative isolation.

What's the status of Federal High-Performance Computing Capabilities?

In 1991, Congress passed the High-Performance Computing Act, establishing an interagency initiative (now called National Information Technology Research and Development (NITRD) programs) and a National Coordination Office for this effort. Eleven agencies or offices participate in the high-end computing elements of the NITRD program. Tables 1a and 1b in Appendix II show the funding level by agency for FY03, the most recent year for which budget data is available. (The overall FY05 budget request for NITRD is $2 billion, but the breakout for the high-performance computing component of that is not yet available.)

The total requested by all 11 agencies in FY03 for high-performance computing was $846.5 million. The largest research and development programs are at NSF, which requested $283.5 million, and the DOE Office of Science, which requested $137.8 million. Other major agency activities (all between $80 and $100 million) are at the National Institutes of Health (NIH), DARPA, the National Aeronautics and Space Administration (NASA), and NNSA. Different agencies concentrate on serving different user communities and on different stages of hardware and software development and application. (Tables 1a and 1b do not include the procurement costs for high-performance computers purchased by agencies, such as NNSA and the National

Oceanic and Atmospheric Administration (NOAA), for computational science related to their missions.[6])

National Science Foundation: In the mid-1980s, NSF established supercomputer centers to serve the academic community. These supercomputing centers provide researchers with access to high-performance computing capabilities and also with the technical support they need to use the facilities effectively. NSF also supports the development of the Extensible Terascale Facility (ETF), a nationwide grid of machines that can be used for advanced communications and data management. The ETF will be coming online in the next year, and a challenge for NSF will be managing the ETF to serve a wide array of users with different scientific computation needs while integrating the ETF with the supercomputing centers.

Department of Energy: DOE has been a major force in advancing high-performance computing for many years. Both the Office of Science and the NNSA invest significantly in high-performance computing. Activities under the Office of Science include the Advanced Scientific Computing Research program, which funds research in applied mathematics, in network and computer sciences, and in advanced computing software tools. In FY04, the Office of Science initiated a new program on next-generation architectures (NGA) for high-performance computing. NNSA uses high-performance computers for simulations and weapons modeling through the Accelerated Strategic Computing Initiative (ASCI).

Defense Advanced Research Projects Agency: DARPA has traditionally focused on hardware development, including research into new architectures. On July 8, 2003, DARPA announced it had selected Cray, IBM, and Sun Microsystems to participate in the second phase of its High-Productivity Computing Systems program. The goal of the program is to provide a new generation of economically viable, high-productivity computing systems for national security and industrial applications by the year 2010.

Other Agencies: NIH, NASA, and NOAA are primarily users of high-performance computing. NIH manages and analyzes biomedical data and models biological processes. NOAA uses simulations for weather forecasting and climate change modeling. NASA relies on high-performance computers for applications including atmospheric modeling, aerodynamic simulations, data analysis and visualization. Scientists at the National Institute of Standards and Technology collaborate with companies and universities to develop high-performance computing applications to address industrial problems. The NSA both develops and uses high-performance computing for a number of applications, including codebreaking. As a user, NSA has a significant impact on the high-performance computing market, but due to the classified nature of its work, the size of its contributions to High-End Computing Infrastructure and Applications and the amount of funding it uses for actual operation of computers is not public.

Interagency Coordination: The National Coordination Office (NCO) coordinates planning, budget, and assessment activities for the NITRD Program through a number of interagency

[6] For example, in FY03 NOAA spent $36 million on supercomputers—$10 million for machines for climate modeling and $26 million for machines for the National Weather Service.

working groups. The NCO reports to OSTP and the National Science and Technology Council. The NCO also manages the HEC-RTF, an interagency effort on the future of U.S. high-performance computing. The HEC-RTF is tasked with the development of a roadmap for the interagency research and development for high-end computing core technologies, a federal high-end computing capacity and accessibility improvement plan, and a discussion of issues relating to federal procurement of high-end computing systems.

7. Witness Questions

The witnesses were asked to address the following questions in their testimony:

Questions for Dr. Marburger

1. What are the Administration's views on the *High-Performance Computing Revitalization Act of 2004*?

2. Please describe the findings and recommendations of the High-End Computing Revitalization Task Force. How will these findings and recommendations be incorporated into the Networking and Information Technology Research and Development program that you oversee?

3. What are the respective roles of the National Science Foundation and the Department of Energy with regard to the provision of high-performance computing resources to university researchers?

Questions for Dr. Wladawsky-Berger

1. How does high-performance computing affect U.S. industrial competitiveness?

2. Are current efforts on the part of the federal civilian science agencies in high-performance computing sufficient to assure U.S. leadership in this area? What should agencies such as the National Science Foundation and the Department of Energy be doing that they are not already doing now?

3. Where are you targeting IBM's high-performance computing research efforts? Are there particular industrial sectors that will benefit in the near term from anticipated high-performance computing developments?

Questions for Dr. Stevens

1. How does high-performance computing affect the international competitiveness of the U.S. scientific enterprise?

2. Are current efforts on the part of the federal civilian science agencies in high-performance computing sufficient to assure U.S. leadership in this area? What should agencies such as the National Science Foundation and the Department of Energy be doing that they are not already doing now?

3. Where should the U.S. be targeting its high-performance computing research efforts? Are there particular industrial sectors or science and engineering disciplines that will benefit in the near term from anticipated high-performance computing developments?

Questions for Dr. Reed

1. How does high-performance computing affect the international competitiveness of the U.S. scientific enterprise?

2. Are current efforts on the part of the federal civilian science agencies in high-performance computing sufficient to assure U.S. leadership in this area? What should agencies such as the National Science Foundation and the Department of Energy be doing that they are not already doing now?

3. Where should the U.S. be targeting its high-performance computing research efforts? Are there particular industrial sectors or science and engineering disciplines that will benefit in the near term from anticipated high-performance computing developments?

APPENDIX I

Section-by-Section Analysis of H.R. 4218, the High-Performance Computing Revitalization Act of 2004

Sec. 1. Short Title

"High-Performance Computing Revitalization Act of 2004."

Sec. 2. Definitions

Amends section 4 of the High-Performance Computing Act of 1991 (HPC Act) to further elaborate on, or amend, the definition of terms used in the Act:

- "Grand Challenge" means a fundamental problem in science or engineering, with broad economic and scientific impact, whose solution will require the application of high-performance computing resources and multidisciplinary teams of researchers
- "high-performance computing" means advanced computing, communications, and information technologies, including supercomputer systems, high-capacity and high-speed networks, special purpose and experimental systems, applications and systems software, and the management of large data sets
- "Program" means the High-Performance Computing Research and Development Program described in section 101
- "Program Component Areas" means the major subject areas under which are grouped related individual projects and activities carried out under the Program

Strikes the definition of "Network" that refers to the National Research and Education Network, which no longer exists as such.

Sec. 3. High-Performance Computing Research and Development Program

Amends section 101 of the HPC Act, which describes the organization and responsibilities of the interagency research and development (R&D) program originally referred to as the National High-Performance Computing Program—and renamed the High-Performance Computing Research and Development Program in this Act. Requires the program to:

- Provide for long-term basic and applied research on high-performance computing
- Provide for research and development on, and demonstration of, technologies to advance the capacity and capabilities of high-performance computing and networking systems
- Provide for sustained access by the research community in the United States to high-performance computing systems that are among the most advanced in the world in terms of performance in solving scientific and engineering problems, including provision for technical support for users of such systems
- Provide for efforts to increase software availability, productivity, capability, security, portability, and reliability
- Provide for high-performance networks, including experimental testbed networks, to enable research and development on, and demonstration of, advanced applications enabled by such networks

- Provide for computational science and engineering research on mathematical modeling and algorithms for applications in all fields of science and engineering
- Provide for the technical support of, and research and development on, high-performance computing systems and software required to address Grand Challenges
- Provide for educating and training additional undergraduate and graduate students in software engineering, computer science, computer and network security, applied mathematics, library and information science, and computational science
- Provide for improving the security of computing and networking systems, including research required to establish security standards and practices for these systems

Requires the Director of the Office of Science and Technology Policy (OSTP) to:

- Establish the goals and priorities for Federal high-performance computing research, development, networking, and other activities
- Establish Program Component Areas that implement the goals established for the Program and identify the Grand Challenges that the Program should address
- Provide for interagency coordination of Federal high-performance computing research, development, networking, and other activities undertaken pursuant to the Program
- Develop and maintain a research, development, and deployment roadmap for the provision of high-performance computing systems for use by the research community in the United States

Leaves substantially unchanged the provisions of the HPC Act requiring the Director of OSTP to:

- Provide an annual report to Congress, along with the annual budget request, describing the implementation of the Program, including current and proposed funding levels and programmatic changes, if any, from the previous year
- Consult with academic, State, and other appropriate groups conducting research on and using high-performance computing

Requires the Director of OSTP to include in his annual report to Congress:

- A detailed description of the Program Component Areas, including a description of any changes in the definition of activities under the Program Component Areas from the previous year, and the reasons for such changes, and a description of Grand Challenges supported under the Program
- An analysis of the extent to which the Program incorporates the recommendations of the Advisory Committee established by the HPC Act—currently referred to as the President's Information Technology Advisory Committee (PITAC)

Requires PITAC to conduct periodic evaluations of the funding, management, coordination, implementation, and activities of the Program, and to report to Congress once every two fiscal years, with the first report due within one year of enactment.

Repeals section 102 of HPC Act, the "National Research and Education Network," which requires the development of a network to link research and educational institutions, government, and industry. This network was developed but has since been supplanted by the Internet.

Repeals section 103 of the HPC Act, "Next Generation Internet," as this program is no longer in existence.

Sec. 4. Agency Activities

Amends section 201 of the HPC Act, which describes the responsibilities of the National Science Foundation (NSF) under the Program. Requires NSF to:

- Support research and development to generate fundamental scientific and technical knowledge with the potential of advancing high-performance computing and networking systems and their applications
- Provide computing and networking infrastructure support to the research community in the United States, including the provision of high-performance computing systems that are among the most advanced in the world in terms of performance in solving scientific and engineering problems, including support for advanced software and applications development, for all science and engineering disciplines
- Support basic research and education in all aspects of high-performance computing and networking

Amends section 202 of the HPC Act, which describes the responsibilities of the National Aeronautics and Space Administration (NASA) under the Program. Requires NASA to conduct basic and applied research in high-performance networking, with emphasis on:

- Computational fluid dynamics, computational thermal dynamics, and computational aerodynamics
- Scientific data dissemination and tools to enable data to be fully analyzed and combined from multiple sources and sensors
- Remote exploration and experimentation
- Tools for collaboration in system design, analysis, and testing

Amends section 203 of the HPC Act, which describes the responsibilities of the Department of Energy (DOE) under the Program. Requires DOE to:

- Conduct and support basic and applied research in high-performance computing and networking to support fundamental research in science and engineering disciplines related to energy applications
- Provide computing and networking infrastructure support, including the provision of high-performance computing systems that are among the most advanced in the world in terms of performance in solving scientific and engineering problems, and including support for advanced software and applications development, for science and engineering disciplines related to energy applications

Amends section 204 of the HPC Act, which describes the responsibilities of the Department of Commerce, including the National Institute of Standards and Technology (NIST) and the National Oceanic and Atmospheric Administration (NOAA), under the Program.

Requires NIST to:

- Conduct basic and applied metrology research needed to support high-performance computing and networking systems
- Develop benchmark tests and standards for high-performance computing and networking systems and software
- Develop and propose voluntary standards and guidelines, and develop measurement techniques and test methods, for the interoperability of high-performance computing systems in networks and for common user interfaces to high-performance computing and networking systems
- Work with industry and others to develop, and facilitate the implementation of, high-performance computing applications to solve science and engineering problems that are relevant to industry

Requires NOAA to conduct basic and applied research in high-performance computing applications, with emphasis on:

- Improving weather forecasting and climate prediction
- Collection, analysis, and dissemination of environmental information
- Development of more accurate models of the ocean-atmosphere system

Amends section 205 of the HPC Act, which describes the responsibilities of the Environmental Protection Agency (EPA) under the Program. Requires EPA to conduct basic and applied research directed toward the advancement and dissemination of computational techniques and software tools with an emphasis on modeling to:

- Develop robust decision support tools
- Predict pollutant transport and their effects on humans and on ecosystems
- Better understand atmospheric dynamics and chemistry

APPENDIX II

Table 1a: Fiscal Year 2003 Budget Requests for High End Computing by Agencies Participating in the National Information Technology Research and Development program (dollars in millions)

Agency	High End Computing: Infrastructure and Applications	High End Computing: Research and Development	Total for High End Computing
NSF	**215.2**	**68.3**	**283.5**
DOE Office of Science	**98.5**	**39.3**	**137.8**
DARPA	16.8	81.9	98.7
NIH	88.2	8.9	97.1
NASA	68.4	26.0	94.4
DOE/NNSA	41.4	39.5	80.9
NSA	--	31.9	31.9
NOAA	13.3	1.8	15.1
NIST	3.5	0.0	3.5
EPA	1.8	0.0	1.8
ODDR&E	--	1.8	1.8
Total:	**547.1**	**299.4**	**846.5**

Source: NITRD National Coordination Office Fiscal Year 2003 Blue Book. The Blue Book is released in August of each year, and thus the data on FY 2003 spending and FY 2004 budget requests levels has not yet been provided to the National Coordination Office.

Note: In addition to the research and development-type activities that are counted for the data included in this table and Table 1b, many agencies devote significant funding to the purchase and operation of high-performance computers that perform these agencies' mission-critical applications.

Acronyms: DARPA—Defense Advanced Research Projects Agency, DOE/NNSA—Department of Energy's National Nuclear Security Administration, EPA—Environmental Protection Agency, NASA—National Aeronautics and Space Administration, NIH—National Institutes of Health, NIST—National Institute of Standards and Technology, NOAA—National Oceanic and Atmospheric Administration, NSA—National Security Agency, NSF—National Science Foundation, ODDR&E—Office of the Director of Defense Research and Engineering.

Table 1b: Funding History from fiscal year 1992 to fiscal year 2003 of high-performance computing research and development programs at various agencies (dollars in millions)

	FY 1992	FY 1993	FY 1994	FY 1995	FY 1996	FY 1997	FY 1998	FY 1999	FY 2000	FY 2001	FY 2002	FY 2003 (Requests)
NSF	**127.00**	**133.90**	**139.10**	**150.00**	**140.32**	**129.20**	**132.90**	**224.70**	**289.80**	**311.70**	**291.50**	**283.50**
DOE Office of Science	**73.00**	**76.20**	**84.60**	**73.10**	**84.49**	**86.00**	**90.53**	**91.90**	**84.10**	**130.30**	**126.70**	**137.80**
DARPA	141.80	169.20	136.20	142.70	77.96	72.70	84.80	48.00	36.50	96.20	81.30	98.70
NIH	8.90	34.40	29.50	29.90	22.40	23.40	23.74	27.10	34.10	59.50	87.20	97.10
NASA	64.00	70.20	84.60	87.40	75.55	88.00	90.10	71.40	124.80	86.80	62.10	94.40
DOE NNSA									113.90	168.30	75.60	80.90
NSA		40.20	32.70	28.20	29.48	30.40	26.42	24.00	31.70	32.90	41.60	31.90
NOAA	1.80	9.40	9.80	2.80	3.30	4.30	4.30	8.80	13.20	12.00	15.60	15.10
NIST	0.90	0.90	0.90	3.60	5.59	4.00	3.99	3.50	3.50	3.50	3.50	3.50
EPA	4.50	6.10	5.90	10.50	8.70	5.60	5.38	4.20	3.90	3.50	1.80	1.80
ODDR&E									2.00	2.00	2.00	1.80
VA					3.00	1.00						
Totals	**421.90**	**540.50**	**523.30**	**528.20**	**450.79**	**444.60**	**462.16**	**503.60**	**737.50**	**906.70**	**788.90**	**846.50**

Source: NITRD National Coordination Office Blue Books, Fiscal Years 1992 to 2003.

Acronyms: DARPA—Defense Advanced Research Projects Agency, DOE/NNSA—Department of Energy's National Nuclear Security Administration, DOE/SC—Department of Energy's Office of Science, EPA—Environmental Protection Agency, NASA—National Aeronautics and Space Administration, NIH—National Institutes of Health, NIST—National Institute of Standards and Technology, NOAA—National Oceanic and Atmospheric Administration, NSA—National Security Agency, NSF—National Science Foundation, ODDR&E—Office of the Director of Defense Research and Engineering, VA—Department of Veterans Affairs.

Program History: Figures from FY 1992-1995 reflect the funding for the High-performance Computing Systems and the Advanced Software Technology and Algorithms Programs. Figures from FY 1996-1999 reflect the funding for the High End Computing and Computation Program. Figures from FY 2000-2003 reflect the funding for the High End Computing Infrastructure and Applications and Research and Development Programs.

Table 2: The top twenty machines of the TOP500 List of the World's Fastest Supercomputers (full list available on line at http://www.top500.org/).

Rank	Site Country/First Year of Operation	Computer / Number of Processors Manufacturer
1	Earth Simulator Center Japan/2002	Earth-Simulator / 5120 NEC
2	Los Alamos National Laboratory United States/2002	ASCI Q - AlphaServer SC45, 1.25 GHz / 8192 HP
3	Virginia Tech United States/2003	X 1100 Dual 2.0 GHz Apple G5/Mellanox Infiniband 4X/Cisco GigE / 2200 Self-made
4	NCSA United States/2003	Tungsten PowerEdge 1750, P4 Xeon 3.06 GHz, Myrinet / 2500 Dell
5	Pacific Northwest National Laboratory United States/2003	Mpp2 Integrity rx2600 Itanium2 1.5 GHz, Quadrics / 1936 HP
6	Los Alamos National Laboratory United States/2003	Lightning Opteron 2 GHz, Myrinet / 2816 Linux Networx
7	Lawrence Livermore National Laboratory United States/2002	MCR Linux Cluster Xeon 2.4 GHz - Quadrics / 2304 Linux Networx/Quadrics
8	Lawrence Livermore National Laboratory United States/2000	ASCI White, SP Power3 375 MHz / 8192 IBM
9	NERSC/LBNL United States/2002	Seaborg SP Power3 375 MHz 16 way / 6656 IBM
10	Lawrence Livermore National Laboratory United States/2003	xSeries Cluster Xeon 2.4 GHz - Quadrics / 1920 IBM/Quadrics
11	National Aerospace Laboratory of Japan Japan/2002	PRIMEPOWER HPC2500 (1.3 GHz) / 2304 Fujitsu
12	Pittsburgh Supercomputing Center United States/2001	AlphaServer SC45, 1 GHz / 3016 HP
13	NCAR (National Center for Atmospheric Research) United States/2003	pSeries 690 Turbo 1.3 GHz / 1600 IBM
14	Chinese Academy of Science China/2003	DeepComp 6800, Itanium2 1.3 GHz, QsNet / 1024 Legend
15	Commissariat a l'Energie Atomique (CEA) France/2001	AlphaServer SC45, 1 GHz / 2560 HP
16	HPCx United Kingdom/2002	pSeries 690 Turbo 1.3GHz / 1280 IBM
17	Forecast Systems Laboratory - NOAA United States/2002	Aspen Systems, Dual Xeon 2.2 GHz - Myrinet2000 / 1536 HPTi
18	Naval Oceanographic Office (NAVOCEANO) United States/2002	pSeries 690 Turbo 1.3GHz / 1184 IBM
19	Government United States/2003	Cray X1 / 252 Cray Inc.
20	Oak Ridge National Laboratory United States/2003	Cray X1 / 252 Cray Inc.